eng

MDPI

Review

Operational Cycle Detection for Mobile Mining Equipment: An Integrative Scoping Review with Narrative Synthesis

Augustin Marks de Chabris *, Markus Timusk and Meng Cheng Lau

School of Engineering and Computer Science, Laurentian University, 935 Ramsey Lake Rd, Sudbury, ON P3E 2C6, Canada; mclau@laurentian.ca (M.C.L.)
* Correspondence: amarks_de_chabris@laurentian.ca

Abstract

Background: Operational cycle detection underpins a range of important tasks, such as predictive maintenance, energy consumption prediction, and energy management for mobile equipment in mining. Yet, no review has investigated the landscape of methods that segment mobile mining vehicle telemetry into discrete operating modes—a task termed operational cycle detection. **Methods**: Following the Preferred Reporting Items for Systematic Reviews and Meta-Analyses, Scoping Review extension (PRISMA-ScR) framework, we searched The Lens database on 27 June 2025, for records published between 2000 and 2025 that apply cycle detection to mobile mining vehicles. After de-duplication and two-stage screening, 20 empirical studies met all criteria (19 diesel, 1 electric-drive). Due to the sparse research involving battery electric vehicles (BEVs) in mining, three articles performing cycle detection on heavy-duty vehicles in a similar operational context to mining are synthesized. **Results**: Early diesel work used single-sensor thresholds, often achieving >90% site-specific accuracy, while recent studies increasingly employ neural networks using multivariate datasets. While the cycle detection research on mining BEVs, even supplemented with additional heavy-duty BEV studies, is sparse, similar approaches are favored. **Conclusions**: Persisting gaps in the literature include the absence of public mining datasets, inconsistent evaluation metrics, and limited real-time validation.

Keywords: operational cycle detection; mobile mining equipment; battery-electric vehicles; machine learning; unsupervised learning

check for updates

Academic Editor: Antonio Gil Bravo

Received: 22 August 2025
Revised: 27 September 2025
Accepted: 29 September 2025
Published: 16 October 2025

Citation: Marks de Chabris, A.; Timusk, M.; Lau, M.C. Operational Cycle Detection for Mobile Mining Equipment: An Integrative Scoping Review with Narrative Synthesis. *Eng* **2025**, *6*, 279. https://doi.org/10.3390/eng6100279

Copyright: © 2025 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (https://creativecommons.org/licenses/by/4.0/).

1. Introduction

Accurate and timely planning of mining operations is paramount to ensuring profitability at the mine site by minimizing equipment downtime. Mobile mining equipment, such as load–haul–dump vehicles (LHDs), haul trucks, and jumbo drills, operate in harsh, unforgiving environments where wear and variability complicate task execution [1]. A key task that allows for more accurate planning is the identification of a mobile mining machine's state of operation at a given time (e.g., identifying that an LHD is scooping up ore). Successful detection of these operating states allows for better predictive maintenance [2], greater control of production [3], and energy optimization [4]. While this task can be described using a variety of terms, we opt to use "operational cycle detection", or "cycle detection". Each of these cycles also consists of discrete operational modes, which we simply refer to as "modes".

Despite the importance of operational cycle detection, there currently exists neither a review nor a system of classification of techniques specifically designed for and applied

to mobile mining vehicles. Recent work reinforces the gap: [5] synthesizes work on heap perception, bucket trajectory planning, autonomous navigation, and monitoring/fault diagnosis, but does not offer a taxonomy for cycle detection on mobile mining vehicles. Likewise, ref. [3] explicitly calls its work one of the first, and [6] catalogues enabling technologies without proposing any cycle detection taxonomy—together underscoring the absence of a focused review. Beyond this absence, the industry's increasing shift to battery-electric fleets (commercial BEV units now operate in more than 20 mines around the world, with more planned [7]) introduces another challenge. Battery electric vehicle (BEV) operational cycles are fundamentally different from diesel ones—in particular, there can be significant variation between the same operational mode performed by similar vehicles, depending on factors like braking frequency. Accurate detection of these nuances is critical: it can help inform battery pack sizing during design, allow State of Charge (SoC) aware dispatch strategies, and optimize charging logistics in daily operations [8].

To address these emerging challenges, this review examines operational cycle detection methods applied to both diesel-powered and battery-electric mobile mining vehicles. The overarching question guiding this review is as follows: which methods have been applied to classify duty cycles in mobile mining equipment, and how do they differ? Framed in Population, Intervention, Comparison, and Outcome (PICO) terms, the population is mobile mining equipment, the intervention is operational cycle detection methodology, the comparator is cross-platform (diesel versus BEV) approaches, and the outcomes are classification accuracy, robustness, and applicability to mining contexts. By mapping this methodological landscape and identifying cross-domain insights, the review aims to close the current knowledge gap and accelerate the development of robust, real-time cycle detection algorithms for next-generation mining fleets, whether they be diesel-powered or electric.

This review closes the current knowledge gap (namely the absence of a Preferred Reporting Items for Systematic Reviews and Meta-Analyses extension for scoping reviews (PRISMA-ScR) guided review of operational cycle detection techniques for mobile mining equipment) by analyzing peer-reviewed studies published between 2000 and 2025. It delivers two contributions:

1. Comprehensive synthesis of operational cycle detection methods for mobile mining equipment.
2. Critical evaluation of persisting research gaps, including data scarcity, inconsistent evaluation practices, and limited real-time validation.

Section 1 covered the necessity for this review and its contributions. Section 2 details the review protocol (databases, inclusion criteria, PRISMA workflow) performed. Section 3 classifies existing diesel operational cycle detection models, and not only surveys the single BEV-specific paper found, but also analyzes three other cycle detection papers involving non-mining BEVs that occur in environments with similar constraints to mining. Section 4 synthesizes the findings, identifies research gaps, and discusses future work, before concluding in Section 5. Through these contributions, the objective of this paper is to accelerate the development of robust diesel and BEV cycle detection algorithms that can be trusted across mines, fleets, and equipment generations.

2. Methods

This review employs a structured scoping review search and selection process for operational cycle detection studies involving mobile mining equipment. No review protocol was prepared or registered for this scoping review.

2.1. Information Sources

On 27 June 2025, we searched The Lens database for records published between 2000 and 2025. No additional bibliographic databases were searched. This database was chosen because (i) its federated index delivers more than 250 million scholarly and 140 million patent records, giving wider grey-literature coverage than Scopus or Web of Science alone; and (ii) it contains the complete Institute of Electrical and Electronics Engineers, Society for Mining, Metallurgy & Exploration and Canadian Institute of Mining, Metallurgy and Petroleum proceedings, where mining automation methods are often first disclosed. In addition, a sentinel set of papers was used for the citation search backward and forward, and to continuously tune the search string. The search string was continuously improved until it reached 90% recall or greater on all records found, and 66% precision or greater for the first 100 records.

2.2. Search

The verbatim Lens query (as run) is given in Supplementary Materials (S1) to enable replication. No automation or query translation tools were used.
Scope boundary. The search was restricted to mobile mining machinery to lower the chance of finding studies that perform cycle detection, but perform them on stationary vehicles, or do not specify whether the vehicle they use is mobile. Because the nomenclature for cycle detection varies widely within the mining industry, we expertly selected terms that reoccurred only in the mining cycle detection literature.

2.3. Eligibility Criteria

Inclusion. Sources of evidence were eligible if they: (i) referred to *mobile* mining equipment; (ii) applied a method to segment sensor data into operational modes; and (iii) used field data or high-fidelity operational logs. *Exclusion.* Simulation-only studies, non-English publications, non-mobile machinery (stationary machines), and studies without operational cycle detection were excluded. *Rationale.* These criteria were chosen to ensure that all included articles would be highly relevant. Focusing on mobile mining machinery excludes adjacent heavy-equipment contexts; requiring segmentation of telemetry into operational modes targets the central concept of cycle detection in this review; field or high-fidelity data were mandated because simulation-only signals rarely capture operational noise and operator variability, reducing ecological validity; and the restriction to only articles in English reflected the capacity of the reviewer's language.

2.4. Selection of Sources of Evidence

Records were screened using the Rayyan [9] application. Rayyan's built-in de-duplication was applied before selection and after merging citation-chasing records. We verified candidates using exact DOI/title matches and title–year–first-author normalization, resolving residual collisions manually. After removing 57 duplicates, 1757 records were processed from the 1814 initially identified (see Figure 1). The selection was carried out in two stages: title/abstracts followed by full text. The title and abstract screening used three sequential yes/no gates: study design, mining machine focus, and presence of an operational cycle detection method. The largest one-stage exclusion occurred at this phase, accounting for 1593 of 1757 screened records (90.7%). Records that passed all gates were sent to full text review, where they were evaluated against every eligibility row. Reasons for exclusion were recorded at both stages, and are tabulated in Supplementary Materials (Table S1) (titles/abstracts) and Supplementary Table S2 (full text). Screening was conducted by a single reviewer; independent duplicate screening and inter-rater agreement

were not performed. Rayyan's automation features were used only to remove duplicates and track records, not for automated exclusion.

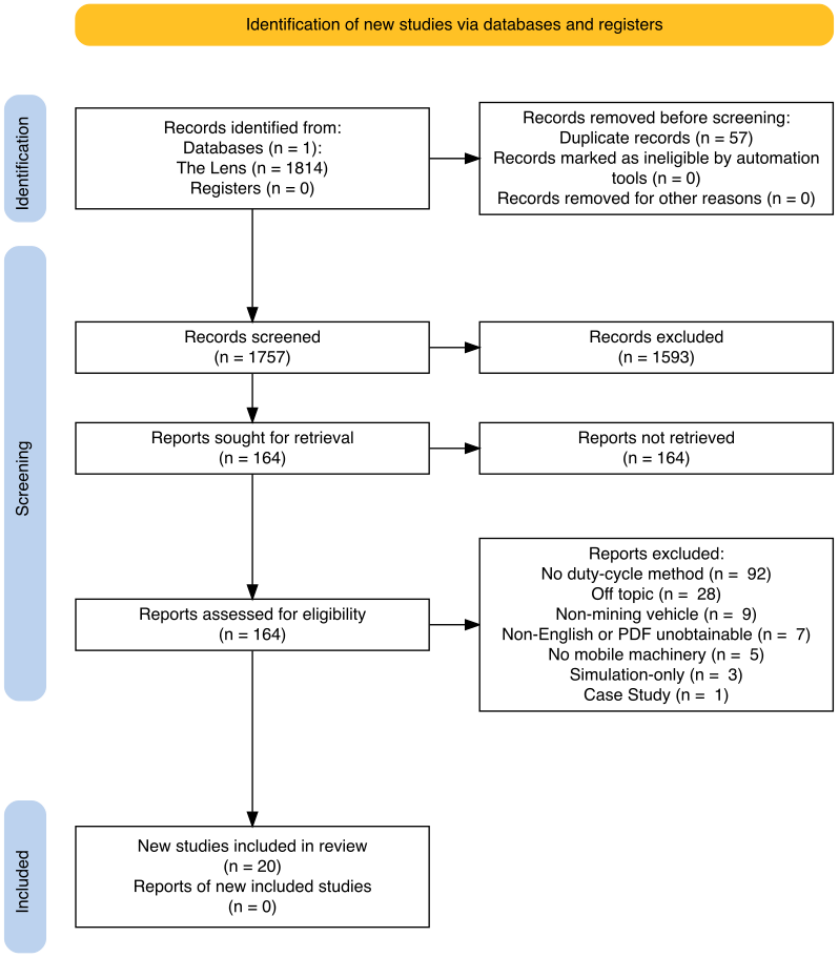


Figure 1. PRISMA-ScR flow diagram illustrating the literature search and screening process.

2.5. Data Charting Process

Data from each included study were charted using a standardized extraction form developed for this review. One reviewer extracted all study characteristics, including vehicle type, mine context, sensors and sampling rate, labelling approach, classification methodology, evaluation metrics, dataset size, and any reported funding or conflicts of interest. No automation tools were used for extraction. No contact with study authors was undertaken, as all required information was available in the published manuscripts.

Data items. Outcomes of interest included classification performance metrics such as per-mode precision, recall, F1-score, overall accuracy, and confusion between adjacent modes, where reported. Other extracted variables comprised study identifiers (authors, year, venue), vehicle class and mining context, sensors and sampling rate, labelling or annotation source, methodological family (rule-based, supervised, or unsupervised), feature engineering steps, dataset size and duration, and any reported funding or conflicts of interest. These items were chosen to allow consistent comparison of methodological approaches across both diesel and BEV studies.

Critical appraisal. Consistent with guidance for scoping reviews, we did not undertake a formal critical appraisal or risk-of-bias assessment of individual studies. Our aim was to map the breadth and diversity of operational cycle detection methods, rather than to evaluate the effectiveness or comparative quality of specific implementations. Instead, methodological features and performance metrics were charted descriptively (Table 1). As

a result, the review does not make quality-weighted judgments on individual studies, but highlights trends, gaps, and methodological commonalities across the evidence base.

Table 1. Per-study results aligned to the review questions.

Study (Author–Year)	Population/Platform	Method and Signals	Data and Validation	Outcomes
Gawelski et al., 2020 [10]	Haul trucks (diesel)	Support Vector Machine (SVM) + DBSCAN on 1 Hz Controller Area Network (CAN) signals (vehicle speed (SPEED), engine revolutions per minute (ENGRPM), selected gear (SELGEAR), brake pressure (BRAKEP), and fuel consumption (FUELS)); 5 s Moving Average; hydraulic oil pressure (HYDOILP) used only for training prep.	∼10 shifts; compared to HYDOILP-based algorithm and manual counts.	Agreement with HYDOILP/operator logs; event-detection reliability ∼ 90%.
Jakobsson et al., 2020 [11]	Mine truck (earthmoving)	CNN vs SVM, Random Forest (RF), K-Nearest Neighbours, Multi-layer Perceptron on 3-axis accelerometer (50 Hz); fixed 2 s windows; auto labels from metadata (payload, speed).	Weeks of data; train 3016/class; oversampled 204,527/class; 80/20 split; Butterworth post-filter.	Balanced test: accuracy > 96%; per-mode True Positive Rate: Idle 0.93, Hauling 0.97, Empty 0.95, Loading 0.95, Unloading 0.94.
Karimi et al., 2021 [4]	Wheel loader (Komatsu WA470)	Hybrid Markov chain + Genetic Algorithm (GA); K-means; smartphone GPS.	∼80,000 points; χ^2 consistency; Simulation vs. EPA/EU cycles.	Constructed cycles consistent; evaluation via fuel/emissions (no classifier accuracy).
Koperska et al., 2020 [12]	LHD (wheel loader)	Convolution of smoothed HYDOILP (1 Hz from 100 Hz) with inverted step; Locally Weighted Scatterplot Smoothing (LOWESS) + variants.	5 working days; >300 cycles; confusion matrix.	Overall acc. 96%; "driving full" sens./prec. ∼98.7%/97.5%; loading sens. 67.6%, prec. 99.5%.
Kozlowski et al., 2019 [13]	Dump truck	Thresholds on ENGRPM+BRAKEP; HYDOILP for verification.	3 test sets (10/40/47 events).	Unloading feasible without HYDOILP; reliability ∼ 90% (limited metrics).
Krot et al., 2020 [14]	Underground haul trucks	Rule-based "virtual sensor" using BRAKEP, SELGEAR, vehicle speed, ENGRPM (1 Hz); HYDOILP ref. 15 Hz for tuning.	3 sets; 97 unloading events; cross-set robustness.	Detection ∼ 90%; false positives ∼ 5%; suitable for online/post.
Markham et al., 2022 [15]	Haul truck (open-pit)	HSMM (unsupervised) on GPS position/velocity (5 s); Viterbi/Expectation Maximization; Fleet Management System for validation.	25 trucks ∼24 h; manual subset of 24 cycles; precision/recall.	Load precision/recall 95.1/98.8%; Dump 87.1/91.1%; Cycle 86.7/90.8%.
Polak et al., 2016 [16]	Loader	Kalman smoothing + statistical break detection on bucket-cylinder pressure (1 Hz).	2 days; 172,800 samples; 207 cycles; manual regime marks.	Accuracy: Unloading 74.4%, Driving full 75.9%, Driving empty 80.2%.
Qi et al., 2023 [17]	LHD (Sandvik LH514)	RF feature selection + Bi-LSTM; baselines VBGMM/SVM/RF/LSTM on multi-sensor (20 params @ 5 s; 5 real-time).	Operational LHD data; weighted/per-mode F1.	Weighted F1 = 91.75%; per-mode F1: Load 0.95, Haul 0.95, Dump 0.83, Transit 0.87.
Saari and Odelius, 2018 [18]	Underground LHD (LH621)	Unsupervised VBGMM on vibration (12.8 kHz) + Cardan speed (5 s).	3 full days; cluster–label "infection" with small manual set; convergence across days.	Reasonable separation of regimes; loading prominent; limited quantitative metrics.
Skoczylas et al., 2023 [3]	Haul trucks; loaders	Deep/conv nets on vehicle speed, SELGEAR, FUELUS, ENGRPM (100 Hz → 1 Hz); trucks partly thresholded.	Initial 17,341 h (trucks)/16,812 h (loaders); downsampled to 2-min/30-s; modified 10 × 10 Cross-validation.	Trucks mean acc. 93% (unloading 96%, loading 90%); loaders mean acc. 58% (one loader ∼ 90%).
Skoczylas et al., 2025 [19]	Haul trucks	VGG16 on IMU (3 accel + 3 gyro, ∼360 Hz); axis autoencoders; rule-based post-process.	1473 cycles; 2462 h; 20 trucks; vs. human detection; data-quality tiers.	With unloading: acc. 82.1%, R 85.8%, P 95.0%; extended (no unload): acc. 87.7%, R 95.7%, P 91.4%.

Table 1. *Cont.*

Study (Author–Year)	Population/Platform	Method and Signals	Data and Validation	Outcomes
Śliwiński et al., 2019 [20]	Haul truck	Autocorrelation function-based cyclicity + rule-based cleaning on 1 Hz multivariate CAN; HYDOILP reference.	Examples over 1 shift and 3 days; visual match.	Suggests vehicle speed + current gear plus one of BRAKE, brake pedal position, fuel consumption, intake air pressure suffice; no numeric accuracy.
Stefaniak et al., 2015 [21]	Mining loader	Kalman smoothing + thresholding on bucket-cylinder pressure.	Not Reported; expert validation.	Identifies "full bucket ride" and "empty backward after unloading"; no numeric metrics.
Timusk et al., 2009 [2]	Electromech. excavator (P&H TS4100); haul truck drives	Supervised Least Squares/Decision Tree/Neural Network/Radial Basis Function/SVM with PCA/Independent Component Analysis/Exhaustive Feature Selection (EFS) on speed-segment and vibration features.	45 h; ∼300 × 1-min records; ∼160 labeled speed profiles; 80% hold-back.	Best KNN+EFS error 6%; empty swing as low as 3%; vibration features improved accuracy.
Wodecki et al., 2018 [22]	LHD (wheel loader)	Compared six smoothing methods on bucket hydraulic pressure to prepare the signal for threshold-based cycle detection (moving average; exponential moving average; linear/robust linear; quadratic/robust quadratic regression).	280 min with 20 known cycles; % cycles detected.	Recovery vs. known 20: MA 105%, EMA 115%, LR 145%, RLR 100%, QR 170%, RQR 100%.
Wodecki et al., 2020 [23]	Drill rigs (FaceMaster 1.7)	Thresholding; KDE-based thresholds; Hilbert demodulation (current); instantaneous frequency (acoustic).	28 holes; compare vs. on-board system.	All 28 holes recovered; regime statistics align across sources.
Wyłomańska and Zimroz, 2014 [24]	Heavy-duty mobile machines (loaders)	Second-moment R & C statistics on ENG_RPM; trapping-event detection.	Simulations + several hours of real ENG_RPM; visual/simulation validation.	R detects traps > 3 samples; C > 2; trapping events 17–28% of signal.
Zhang et al., 2019 [25]	Electric-drive mining truck (220 t)	PCA + density-peak clustering on speed/power/torque/grade; 10 s "short-stroke" windows.	146 kinematic sequences; representative cycle vs. field data.	Three categories (loaded/high power; idle/no power; braked/high speed); representative mirrors field.
Stachowiak et al., 2020 [26]	LHD (wheel loader; hydraulic)	HYDOILP 1 Hz; ten smoothing methods: MA, Weighted Moving Average (WMA), Exponential Moving Average (EMA), Double-Exponential Moving Average (DEMA), Hull Moving Average (HMA), Arnaud Legoux Moving Average (ALMA), Linear Regression, Quadratic Regression, LOWESS, Kalman; cycle detection via convolution with inverted step (rule thresholds).	3 h total across two samples: 1 h with 11 actual cycles; 2 h with 15 actual cycles; MA window 70 samples; validation by difference from actual counts and duration statistics.	LOWESS best: correct count for both samples with minimal shift. Sample 1 (11): most methods 0 diff; HMA +1; LR −1; Kalman −2. Sample 2 (15): LOWESS 0; MA +1; EMA +5; DEMA +13; WMA +5; HMA +10; ALMA +2; LR 0; QR +3; Kalman +2. Kalman delayed; DEMA/HMA produced spurious short cycles.

Across the corpus, the choice of the model follows (i) which signals are available and at what rate, and (ii) the density of ground-truth labels. Low-rate supervisory streams (10 Hz; speed/gear/torque, GPS 0.2–1 Hz) with sparse or batch labels align with transparent rules, Hidden Markov Models (HMM), Hidden Semi-Markov Models, or clustering; with dense or reliably generated labels, the same streams support temporal deep models (CNN/LSTM/BiLSTM) that capture multi-minute cycle structure. Mid/high-rate sensors (50 Hz; hydraulics, IMU) without labels align with change-point detection and probabilistic clustering; with expert labels, convolutional networks learn features directly from windows. For BEVs, power/SoC/DC-bus signals (1–10 Hz) make energy-aware sequence models attractive (regen/grade effects), while edge-latency constraints often motivate hybrids (rules + lightweight HMM or tiny CNN). This mapping underpins the comparative synthesis in Section 4.

2.6. Synthesis of Results

Extracted data were synthesized descriptively, consistent with scoping review methodology. To facilitate comparison, studies were grouped along two axes: equipment type (diesel versus battery-electric) and methodological family (supervised or unsupervised/semi-supervised). Within these groups, we charted study characteristics (Table 1) and reported outcomes in their original units (e.g., accuracy, precision/recall, F1-scores, energy- or cycle-time errors). Given the heterogeneity of metrics and contexts, no meta-analysis or formal critical appraisal was undertaken. Instead, trends, gaps, and methodological commonalities were summarized in a narrative style.

2.7. Hand-Selected Exemplars: Rationale and Use

Although only one peer-reviewed paper addresses operational cycle detection for underground BEVs, several heavy-duty electric platforms have begun to tackle analogous challenges (e.g., high payload variation, steep gradients, and intense auxiliary loads). To reinforce Section 3.3 and broaden its relevance, we include three additional studies drawn from adjacent BEV domains: a construction loader, an industrial forklift, and an articulated transit bus. These studies were selected because (i) their operational profiles exhibit key parallels with mobile mining machinery, such as repeated load-transport–dump-transit cycles, mixed-mode traction demand, stop-and-go operation, and energy recovery via regenerative braking; (ii) report end-to-end pipelines with transparent data engineering and feature construction; and (iii) articulate evaluation metrics clearly enough to support qualitative comparison.

3. Results

3.1. Study Selection and Corpus Characteristics

The structured search in The Lens retrieved 1567 records; backward and forward citation chasing of a predefined sentinel set added 247, bringing the total to 1,814 records (see Supplementary S1 for the full Boolean string). The final search string recovered 11 of 12 sentinel papers, achieving 92% recall while maintaining 66% precision over the first 100 records. A single sentinel set paper (Lewis, 2018 [27]) was not captured by either method and is cited only for context. After screening, 1593 records were excluded, leaving 164 articles for full-text review. Of these, 144 were removed for being simulation-only, non-mining, or lacking a segmentation method. Twenty empirical studies met all eligibility criteria (19 diesel, 1 electric-drive) and were included in the synthesis. The PRISMA flow diagram (Figure 1) illustrates the selection process, and full source details are provided in Tables S1 and S2.

Across the 20 included studies (19 diesel, 1 electric-drive), extracted study characteristics are summarized in Table 1. These include vehicle class, mining context, sensors and sampling rate, labelling approach, methodological family, dataset size, performance outcomes, and reported funding. The following subsections organize the results thematically by methodological family. In addition, a single article from the sentinel set (Lewis [27]), not captured by the formal search, is occasionally referenced to provide methodological context. This study is not part of the included corpus and is not reflected in the PRISMA counts or tables.

3.2. Operational Cycle Detection-Diesel

3.2.1. Sensors, Data Processing and Feature Engineering

Cycle-detection models rely on access to data. Mobile mining equipment captures these data through a variety of sensors onboard, with four main 'families' dominating current diesel-based approaches. Each family provides distinct and physically meaningful signals that help distinguish between operational cycle modes.

- Hydraulic system pressure: These sensors measure pressure fluctuations in the lift, tilt, and pump circuits of the hydraulic system. Their near-rectangular pulses clearly segment the loader cycle into phases such as loading, hauling, dumping, and transiting.
- Driveline kinematics: This family includes gear selection, engine revolutions per minute (RPM), and ground speed sensors. These signals help infer machine intent and movement patterns, such as identifying dumping while stationary through a spike in RPM at zero speed.
- Brake-system hydraulics: These sensors capture pressure changes in the service and retarder brake lines. The spikes in these signals serve as strong indicators of dumping or controlled descents on ramps, which aid in phase detection.
- Auxiliary streams: This group comprises lower-rate signals, such as fuel flow and turbo boost, and higher-rate streams, such as those originating from Inertial Measurement Units (IMUs) and the Global Positioning System (GPS). These enrich cycle detection models by providing load, terrain, and location cues; especially useful when primary sensors are unavailable or noisy.

A common challenge is the mismatch in sensor sampling rates, which can complicate downstream data analysis. To ensure consistency, high-frequency sensor data are typically resampled and stored on a unified timeline. This processing step allows for standardized downstream processing, such as sliding-window segmentation for Convolutional Neural Networks (CNNs), where segment lengths can range from 30 to 120 s [3]. In some high-rate vibration studies, such as [18], even more extreme sampling rates (12.8 kilohertz) are down-sampled to align with slower sensor streams, allowing simultaneous analysis of disparate streams.

Given the noisy and highly variable conditions under which sensor data are collected, many studies introduce a data preprocessing step to mitigate noise and inconsistencies. This can involve temporal smoothing (reducing high-frequency fluctuations while preserving underlying trends), using methods such as Locally Weighted Scatter Plot Smoothing (LOWESS), Kalman filtering, or robust linear and quadratic regression [10,16,22]. In addition to smoothing, data processing can include statistical outlier removal, where values outside of a statistically defined range are removed to improve prediction accuracy [18,19]. Variables are frequently normalized via min-max scaling or expressed as proportions of their maximum observed values to enhance comparability between samples [14]. Feature selection techniques are then used to reduce dimensionality, including autocorrelation-based similarity filtering [20], exhaustive combinatorial search [2], or ranking by importance with a Random Forest model [17].

Once predictions are made, a final step is often employed to suppress label flips; rapid, isolated changes in predicted mode that typically arise from momentary noise or classifier uncertainty, such as a single misclassified sample between two stretches of consistent labels (e.g., five seconds of loading, a half-second of dumping, followed by five seconds of loading). These include majority voting over sliding windows and finite-state filtering approaches that impose temporal consistency by modeling state transitions with memory [19,27].

High-rate sensors (50 Hz)—accelerometers/IMUs from tens of hertz up to the kilohertz range and hydraulic pressures sampled at tens to hundreds of hertz—are frequently exploited without detailed labels: change-point methods and probabilistic clustering (e.g., variational Bayesian GMMs) discover boundaries and modes from engineered features. When expert labels are available, convolutional networks operating on raw or minimally processed windows can learn discriminative features directly; where edge latency or bandwidth is constrained, autoencoders compress high-rate signals before classification, or lighter models (e.g., trees/Random Forests) replace heavier CNNs for order-of-magnitude speed-ups. In fault-tolerant deployments, simple rule pipelines over robust CAN signals remain valuable as a fall-back when more informative but failure-prone channels (e.g., auxiliary pressure transducers) degrade.

3.2.2. Supervised Methods

Supervised cycle detection models begin with a training set in which every slice of sensor data has been tagged as a mode (e.g., "loading", "hauling", "dumping"). By examining how signals behave during each tagged period, the model aims to distinguish one activity from another.

Early operational cycle detection systems relied on single-signal thresholds: if a measurement crossed a preset limit, the system declared a mode change. Although conceptually elegant and computationally frugal, these heuristic-based models falter when faced with vibration noise, sensor failure, and ambiguous stops.

Polak [16] pioneered this approach with a two-stage routine tailored to hydraulic-bucket pressure data sampled at 1 Hertz (Hz). First, the raw pressure trace was smoothed using a Kalman filter to suppress noise while preserving the signal edges. After filtering, a 6 MPa threshold was applied to distinguish between full and empty bucket states. Subsequently, a variance drop detector segmented the timeline into loading, haulage, and return phases. Despite relying on a single sensor channel, this method achieved a classification accuracy of 75%.

Wodecki et al. [23] identified noise suppression as the primary performance bottleneck. Of the methods tested, only robust linear and quadratic regressions applied over a 61-s window were able to smooth a hydraulic pressure trace into a near square-wave form, successfully recovering all 20 ground-truth cycles.

Building on this insight, Stachowiak et al. [26] benchmarked ten smoothers on HYDOILP (1 Hz) with a fixed downstream detector (LOWESS/MA-family/LR/QR/ALMA/DEMA/HMA/Kalman → inverted-step convolution + extrema), showing LOWESS most faithfully preserved cycle structure across irregular/regular runs; LR was a viable second, while DEMA/HMA over-segmented around loading transients and Kalman lagged/over-smoothed short cycles—only LOWESS and LR recovered the correct 15 cycles.

In parallel, Koperska et al., Koperska et al. [12] introduced a two-stage HYDOILP (1 Hz) pipeline—LOWESS + inverted-step convolution to bracket "drive full/empty," followed by a local smooth-and-convolve to mark loading—validated on >300 cycles over five shifts: full/empty reached >96% accuracy with >97% precision/recall, while loading

was high-precision (99.5%) but low-recall (67%), making it a robust boundary detector and a reliable label source for downstream ML.

Single-sensor systems can fail when the sensor is damaged or when multiple stoppage-like states look identical on that channel [14]. Researchers, therefore, began to merge readily available variables into more comprehensive data sets.

Krot et al. [14] compensated for a failed hydraulic pressure sensor in haul trucks by implementing a set of four signal rules. A two-stage pulse filter, applied to brake pressure, engine RPM, gear position, and travel speed, successfully identified 90% of unloading events with only 5% false positives, while also delineating complete cycle boundaries.

Śliwiński et al. [20] analyzed 19 1 Hz sensor streams, first identifying those whose temporal patterns correlated with the hydraulic-oil-pressure channel. They then filtered to stationary periods (speed = 0, gear = 0) where fuel consumption exceeded 10% of peak. Using only three common signals—speed, gear, and throttle position—they were able to reconstruct full work cycles even when hydraulic pressure data were unavailable.

Wodecki et al. [23] adapted rule-based segmentation to drilling rigs, applying kernel density thresholds to 10 Hz current-draw data to distinguish idle and drilling modes. In parallel, a Hilbert envelope transform of 48 kHz acoustic data was used to verify the percussion frequency. This dual-signal approach successfully segmented a full shift into 28 boreholes, and even flagged instances of incorrect drilling.

In these studies, combinations of variables (for example, 'brake engaged + neutral gear + RPM spike = unload') act as anchor points. These anchors are then expanded into dense, sample-level annotations using simple temporal logic such as pulse counting, windowed autocorrelation function similarity, or regime-specific thresholds. The result is a fault-tolerant, light calibration pipeline capable of both online deployment and retrofitting historical datasets without manual labelling. However, even multivariate rule sets falter when operator habits or machine variants break the underlying hand-coded assumptions: Limitations that motivate a transition towards data-driven models.

Unlike rule-based models, data-driven approaches replace hand-tuned thresholds with algorithms that learn the mapping from sensor streams to operational states; in effect, building their own set of rules. Their evolution mirrors the richness of the available data.

Timusk et al. [2] cast loaded vs unloaded-swing detection on an excavator as a supervised task. Forty-five hours of speed–vibration data were manually labelled; time-domain and statistical features were exhaustively pruned, and even elementary classifiers (k-Nearest Neighbors, Linear Discriminant Analysis, Support Vector Machine (SVM)) reached 94% accuracy (6% error)—proof that judicious features let simple data-driven models rival rule-based heuristics.

Jakobsson et al. [11] fed a one-dimensional CNN with 50-Hz accelerometer windows and labels derived automatically from the mass and speed of the payload. The model hit 96% accuracy, beating classic baselines without feature crafting.

Skoczylas et al. [19] adopted a fully off-board approach in which six axis-specific auto-encoders compressed 360 Hz inertial IMU windows into low-dimensional codes. These representations were then passed to a one-dimensional Visual Geometry Group 16-layer network, a convolutional neural network (CNN) architecture widely regarded as one of the most influential and high-performing models in computer vision [28]. Applying logical sequence rules to the raw outputs improved classification performance, yielding an accuracy of 87.7%, with 91% precision and 96% recall.

Lewis [27] trained a compact CNN on 10 Hz telemetry from a Caterpillar LHD; camera-derived labels plus morphological and Markov smoothing lifted accuracy from 66% to 80%.

Qi et al. [17] paired Random-Forest Feature Selection with a Bi-LSTM network on five-second Controller Area Network (CAN) snapshots from a Sandvik LHD, achieving 91.8% weighted accuracy.

Skoczylas et al. [3] auto-searched neural network topologies that used a variety of sensor streams; hydraulic pressure peaks served as weak labels. The best CNN scored 93% unloading accuracy for haul trucks but only 57% for LHDs.

3.2.3. Unsupervised and Semi-Supervised Methods

A major challenge with supervised cycle detection methods is the requirement for every data point to be labelled in advance (e.g., 'loading', 'hauling'). Creating these labels is slow and difficult. Experts usually need to carefully examine sensor data, cross-check them with operational logs, and align the signals with specific events. In underground mines, this process is even harder due to rough working conditions, faulty sensors, and constant changes in the way machines operate. These issues often make reliable manual labelling unrealistic. To address these challenges, researchers have turned to alternative approaches that reduce or eliminate this burden.

Unsupervised learning methods aim to infer structure in the data without relying on preexisting labels, typically by grouping similar data points (also known as clustering). Semi-supervised learning offers a middle ground between unsupervised and supervised learning by leveraging a small amount of labelled data alongside a larger pool of unlabeled samples. In mining applications, this can translate to significant annotation cost savings while still maintaining classification performance.

Wyłomańska and Zimroz [24] pioneered a label-free approach to cycle detection on diesel LHDs by monitoring a single channel: engine RPM. They computed two complementary second-moment statistics—windowed variance (R) and cumulative energy (C)—to flag "trapping events" (stationary periods) whenever both metrics fell below an empirical 50 RPM threshold. Used individually, R tended to miss extended idle periods, while C tended to over-flag short stalling events. In combination, however, they identified 95 robust events (approximately 18% of the total shift), which corresponded almost exclusively to either idling at 700 RPM or sustained full-throttle operation at 2000 RPM.

While Wyłomańska and Zimroz [24] demonstrated that careful signal analytics can eliminate the need for manual labelling, the study also revealed two key limitations: (i) thresholds require manual tuning for each machine, and (ii) subtle or multi-phase activities become indistinguishable when relying on a single variable. To address these constraints, subsequent research has integrated unsupervised or weakly supervised machine learning with domain expertise, aiming to balance automation with interpretability.

Markham et al. [15] illustrate the most label-frugal end of the spectrum by converting five-second GPS data from 25 haul trucks into a ten-symbol sequence that encodes vehicle speed and proximity to shovels, stockpiles, and crushers. Without supplying a single hand annotation, they fit a Hidden Semi-Markov Model (HSMM) whose transition matrix was constrained only by obvious physical rules (e.g., a truck cannot jump directly from loading to empty transit). Compared to the mine's curated fleet management log, the HSMM recovered 99% of loading events and 91% of dumps, even uncovering 24 valid operational cycles that the official record had missed.

Rather than eliminating labels entirely, Gawelski et al. [10] focus on identifying one mode (unloading) with high confidence and letting that cue propagate through the data set. A compact Support Vector Machine, trained on just two labelled shifts, recognized unloading whenever vehicle speed was zero and engine RPM sat between 1200 and 2500 RPM. These unloading data points are clustered with the Density-Based Spatial Clustering of Applications with Noise (DBSCAN) algorithm to identify successive cycle boundaries;

everything between two clustered unloads was treated as a full operational cycle. The result was a complete cycle trace without consulting hydraulic or brake pressure channels, which are often noisy or missing in legacy data acquisition systems.

Saari and Odelius [18] inverted the previous workflow: they allowed the algorithm to discover latent patterns first, then inject semantic meaning. A variational Bayesian Gaussian mixture model (VBGMM) ingested millions of 12.8 kHz vibration samples, augmented with speed, and clustered them into ten groups entirely label-free. Subsequently, less than 1% of the samples received expert-assigned mode labels; the group that captured most of a label inherited that operational label. With only a single minute's worth of labelled data, the VBGMM achieved 100% accuracy in identifying the idling mode and produced clean, easily interpretable partitions for loading, hauling, and transiting modes.

An overview of the major semi-supervised and unsupervised approaches is covered in Table 2.

Table 2. Label-efficient learning strategies for operational cycle detection.

Strategy	Key Study and Method	Label Economy	Headline Result
Fully generative discovery	Markham et al. (2022) [15]: Hidden Semi-Markov Model on symbolized five-second GPS fixes; physics rules only prune impossible transitions.	0 manual labels	Retrieved 99% loading, 91% dumping modes; uncovered 24 cycles missing from the fleet log.
Event-anchored seeding	Gawelski et al. (2020) [10]: SVM learns the *unloading* pattern; DBSCAN clusters unload flags to bracket cycles.	Labels for one sub-event	Full duty-cycle trace generated without hydraulic or brake-pressure channels.
Cluster-first infection	Saari and Odelius (2018) [18]: VBGMM finds ten vibration-speed clusters; $<1\%$ expert tags "infect" clusters with regime labels.	Minutes of tags	Identified idling at 100% accuracy; clean partitions for loading, hauling, transit.

3.2.4. Diesel Specific Trends and Key Limitations

Early cycle detection work on diesel haulage equipment relied almost exclusively on single channels (typically hydraulic cylinder pressure) mixed with fixed thresholds or simple statistical rules [16,24]. Over the past decade, researchers have migrated toward multivariate inputs and increasingly sophisticated learning algorithms.

"Virtual-sensor" rule sets now fuse standard electronic control unit (ECU) signals such as brake pressure, gear, engine speed, and travel speed to withstand missing hydraulic pressure sensors [10,14], while shallow classifiers demonstrated as far back as 2009 that well-curated features enable reliable mode labelling without hand-tuned thresholds [2]. The latest generation of papers adopts deep networks—CNNs on vibration streams [3,11] and Bi-LSTMs on five-second CAN snapshots [17] to capture non-linear, context-dependent dynamics. However, only 4 of the 19 diesel studies run rule-based and data-driven models side by side on a common dataset, leaving practitioners without a benchmarked answer to the question 'Which approach is best for my fleet?'

Another consistent finding is the strong imprint of operator style on sensor signatures. Cycle duration, braking patterns, and even average RPMs per mode change with individual driving habits or changing haul routes, can erode model accuracy [10,27]. Few papers attempt domain adaptation or style-invariant training; most simply caution that larger and more diverse datasets are needed.

The widespread fragility of hydraulic pressure sensors (frequently damaged or completely absent) has made many legacy rule-based methods inoperable, prompting a shift to more reliable ECU-derived signals such as brake pressure, engine speed, and gear selection [12,20].

Finally, the labor-intensive nature of supervised learning remains a significant bottleneck: high-performing models often require extensive manual labeling, such as annotating 45 h of excavator operation [2] or frame-by-frame video analysis of underground LHDs [27]. This annotation burden scales poorly and impedes fleet-wide deployment.

3.2.5. LHD vs. Truck Transferability

Evidence consistently indicates that cycle detection models developed for haul trucks do not transfer well to LHD machines. Using identical non-hydraulic inputs (e.g., speed, gear, fuel rate), a CNN reached approximately 93% accuracy on haul trucks but only 58% on LHDs, with systematic confusion between loading and dumping operations and marked instability across machines [3]. By contrast, loader-specific fusion and sequence models have achieved substantially higher performance: a Random Forest feature selection combined with a Bi-LSTM reported a weighted F1 of 91.8%, with loading at 95.4% and dumping at 83.1% [17].

Hydraulic-based approaches highlight the same limitation. While LOWESS smoothing with change-point detection yielded 96% accuracy for "driving full/empty bucket" states, accuracy for loading was only 67%, with systematic errors at shift boundaries [12]. Similar hybrid models (those that combine different algorithms) reported moderate correctness for driving states (76–80%) but could not differentiate between loading and dumping [16,26].

These findings suggest that LHD duty cycles pose greater algorithmic challenges than trucks due to shorter and more variable phases, the inclusion of digging operations, and the heavy imprint of operator style. The most informative signal, hydraulic pressure, is simultaneously noisy, vibration-contaminated, and prone to sensor failure, while non-hydraulic cues alone fail to distinguish loading from unloading. Consequently, cycle detection algorithms designed for trucks require refinement in order to detect the cycles of loaders.

3.3. Operational Cycle Detection-Battery-Electric

No peer-reviewed study specific to mining BEVs met the eligibility criteria; we therefore include one mining electric-drive (diesel-electric) haul-truck cycle study as adjacent evidence. Its characteristics (vehicle type, mine setting, sensor sources, method, data set size, and results) are included in Table 1. Because no formal critical evaluation was performed, the findings of this study are descriptively reported alongside three heavy-duty BEV analogs to contextualize the methodological approaches and highlight transferability considerations.

3.3.1. Unique Challenges of Battery-Electric Vehicle Operational Cycle Detection

Battery electric vehicles (BEVs) exhibit an operational cycle behavior that differs fundamentally from that of internal combustion engine vehicles. Unlike diesel vehicles, BEVs feature the following:

- **Aggressive speed transients**: Instantaneous peak torque enables sharper accelerations and decelerations. BEV acceleration phases can be up to 26% longer, and with 18% higher magnitude, than comparable ICE phases—leaving less time for steady cruising and complicating phase detection [29].
- **Regenerative-braking distortion**: Deceleration profiles are structurally different due to regenerative braking. These introduce negative power spikes and longer braking windows that alter the temporal features classifiers must learn to recognize [8,29].
- **Load asymmetry and State of Charge (SoC) variability**: BEVs suffer from power derating under low battery or high temperature, making the same task look different in different contexts (a major challenge for supervised models that assume consistency).

Given these challenges, BEV cycle-detection algorithms must not only infer vehicle activity but also deal with spikes in SoC due to battery swapping, and distinguish rising SoC due to charging versus regenerative braking. Mining-specific research in the BEV domain remains extremely sparse, with only one dedicated study identified through the database search.

3.3.2. Electric Haul-Truck Cycle Detection

Zhang et al. [25] report the only retrieved study attempting to construct a representative driving cycle for an electric-drive mining truck from onboard telemetry. The platform, a 220 ton haul truck operating on a 3.5 km open–pit route with two short ramps, one long incline, and compacted gravel (estimated 3% rolling resistance), completed six full load–haul–return cycles plus an idle transfer. These runs were segmented into 146 fixed 10-s analysis windows (approximately 1460 s total), each described by ten kinematic and energy-relevant statistics, including speed, traction-power ratio, and average road grade.

The resulting feature matrix was standardized and reduced with Principal Component Analysis (PCA). In this lower-dimensional space, Density-Peak Clustering distinguished three behavioral groups: high-power loaded hauling (C1), low-power idling or coasting (C2), and high-speed dynamic braking return runs (C3). Importantly, the 10-s windows served as analysis units for clustering, not as the final cycle itself. To produce a realistic operational template, the authors subsequently assembled a synthetic 1500-s cycle by scaling representative windows from each cluster to preserve the observed loaded-haul/empty-return distance ratio. The final profile comprised 930 s of loaded haul, 170 s idle, and 400 s regenerative return, corresponding to 3.4 km loaded and 3.25 km empty. The reconstructed cycle's speed, power, and grade traces closely mirrored the original telemetry, which the authors deemed "strongly reflective" of real operation.

The study demonstrates how clustering-based cycle construction can reproduce field statistics with high fidelity, but its scope remains narrow: one vehicle, a single route, and fixed window lengths. Moreover, the platform is described as an "electric-drive" truck, which in mining contexts often refers to diesel-electric haulage rather than a true battery-electric vehicle. Clarifying this distinction is essential for correctly situating the study within BEV duty-cycle research. In either case, the work highlights both the promise of unsupervised pipelines for heavy-duty vehicles and the immaturity of the mining BEV cycle-detection literature compared to its diesel counterpart.

3.3.3. Adjacent Heavy-Duty BEV Studies

Although only one peer-reviewed study focusing specifically on underground BEVs was identified, this review also examines three additional operational cycle detection studies from BEVs that are used in mining for non-production tasks—loaders and forklifts are commonly used in both surface and underground operations, and buses may be used to shuttle personnel. These articles are included not to prescribe ready-made solutions, but to illustrate how cycle detection techniques have been applied in contexts that share key operational characteristics with mining BEVs, such as heavy-duty workloads, terrain variability, and prevalent regenerative braking.

Several heavy-duty electric cycle detection articles have begun to tackle analogous challenges: high payload variation, steep gradients, and intense auxiliary loads. The three studies are summarized in Table 3.

Table 3. Heavy-duty BEV analogs and their rationale for inclusion.

Article	Vehicle Studied	Reason for Inclusion
Ren et al. 2024 [30]	5 t electric wheel loader	Loader-centric segmentation that embeds variable-mass and shovel-resistance peaks; demonstrates a CAN-bus → PCA + K-means pipeline with <3.5% energy-error fidelity.
Tong & Guan 2024 [31]	Counter-balanced battery-electric forklifts	Double-layer Markov model captures dynamic payload profiles; large-scale (194 M points) micro-trip synthesis keeps 13 performance metrics within ±1%.
Tong & Ng 2023 [32]	Battery-electric buses on steep Hong Kong routes	Gradient-aware micro-trip assembly reproduces speed–acceleration statistics on ±6% slopes—an analog for deep-ramp haulage.

Percentages denote absolute error relative to measured values.

Ren et al. [30] begin with the classic "V-pattern" of construction loader work. After cleaning 20 CAN-logged cycles, they segment data by gear, speed, and hydraulic–pressure gradients, compress ten variables into three principal components, and then cluster 200 segments with K-means. Representative segments nearest each centroid are concatenated and scaled, yielding a synthetic cycle whose simulated energy demand differs from field tests by only 0.01 kWh—proof that mass variation and shovel-load peaks must be hard-wired into BEV loader cycles.

Tong and Guan [31] address the challenge of a forklift's constantly changing cargo weight. They process three months of 50 Hz telemetry by filtering it into 11,595 micro-trips, applying PCA for dimensionality reduction, and clustering the results. These clusters are then sequenced using a double-layer Markov chain—an 18 × 18 transition-probability matrix (TPM) at the macro-state level, paired with velocity–acceleration TPMs at the micro-state level. A genetic algorithm further refines the sequence until 13 statistical performance metrics match the ground truth within 1%.

Tong and Ng [32] extend the micro-trip idea to electric buses on steep urban routes. GPS data are grouped into downhill, flat, and uphill modes; gradient-sensitive vehicle-specific power is computed; and sub-cycles are concatenated until weekday and weekend profiles match 13 speed acceleration descriptors within 10%. The resulting Electric-Bus Driving Cycles with Road Gradient (EBDCRG) expose how continuous 6% slopes dampen aggressive driver inputs, paralleling the ramp effects faced by underground BEVs.

Crucially, each study isolates one mining-relevant complication: mass change, payload change, or steep gradients. Cycle detection systems for mining BEVs, particularly underground mining BEVs, must accommodate all three simultaneously. The techniques above provide a starting toolkit, but a comprehensive mining solution will require fusing mass, gradient, and payload-aware features.

Applicability considerations. Adjacent BEV platforms are included (e-wheel loaders, forklifts, e-buses) strictly to compare methods, not environments: IMU/CAN-based segmentation, micro-trip clustering, grade- and payload-aware features, and sequence models (CNN/LSTM/HSMM) are generally transferable. Underground conditions—GNSS denial, dust and vibration, confined headings, and steep gradients—limit transferability; therefore, GNSS-based features and vision methods sensitive to environmental degradation are treated as non-transferable in this review.

3.3.4. Emerging Trends and Limitations

Although only four heavy-duty BEV studies are reviewed, a converging workflow emerges. All reject passenger vehicle operational cycles in favor of vehicle and context-specific templates that capture the dominant mechanical driver of energy use (whether it is shovel resistance, cargo mass, road slope, or haul/return asymmetry). Continuous telemetry is first chopped into short kinematic strokes or micro-trips, then passed through

PCA, which typically compresses a dozen raw variables into three axes explaining over 80% of the variance. Clustering (usually K-means, though [25] employs Density–Peak and [31] fuse a double-layer Markov model with a genetic algorithm) labels these segments, after which representative windows are concatenated into a synthetic cycle. Fidelity is judged quantitatively: loader and forklift cycles reconstructed in a simulated environment only deviate <5% from measured energy use [30,31], setting an informal benchmark for future work.

The same papers also expose persistent limitations of current BEV operational cycle detection techniques. Each cycle is tightly bound to a single machine, operational cycle, or route (e.g., one Hong Kong bus line, one V-pattern loading task, one representative haul), so transferability remains untested. Dataset sizes are modest, raising questions about statistical representativeness. Selective simplifications creep in, such as omitting lift-lower episodes in forklifts or passenger load in buses, potentially biasing energy estimates. Validation practices are also heterogeneous: while most report error rates, the haul-truck study by Zhang et al. [25] relies on qualitative "good agreement", underscoring the need for standard metrics across platforms.

4. Discussion

This review addressed the overarching question of which methods have been applied to detect operational cycles in mobile mining equipment and how methodological assumptions differ. In relation to the objectives set out in Section 1, the evidence synthesis shows that: (i) diesel studies have evolved from rule-based thresholds to multivariate deep learning architectures, (ii) the nascent BEV literature—both mining and heavy duty analogs favor unsupervised clustering pipelines with PCA-based compression, and (iii) significant gaps remain around open datasets, evaluation standards, and readiness for deployment. These findings inform both methodological practice and future research directions.

4.1. Trends in Operational Cycle Detection

Four overarching trends have been identified in operational cycle detection methodologies.

1. **Shift from Thresholds to Learned Representations**: Across both diesel and BEV platforms, there is a clear progression from hand-crafted rule sets to learned representations. Diesel studies have moved from single-variable and threshold-based to supervised deep learning architectures such as CNNs and Bi-LSTMs [3,17]. Similarly, BEV studies employ PCA-based feature compression and clustering pipelines that allow for flexible, data-driven phase discovery without predefined thresholds [30,31].
2. **Emergence of Modular, Pipeline-Oriented Architectures**: A common architectural pattern is becoming apparent: segmentation, then dimensionality reduction, then clustering or classification, ending with post-processing or synthesis. This modular structure is evident in both unsupervised and supervised workflows, including [20] for diesel and [25,30] for BEVs. It offers a flexible foundation for adapting classifiers across platforms and signal types.
3. **Toward Quantitative, Simulation-Based Validation**: Earlier work often reported only qualitative agreement or raw cycle counts. Recent studies are trending toward more standardized, simulation-based validation: diesel classifiers are now evaluated using segment-level precision and recall [17,27], while BEV studies benchmark reconstructed cycles using energy-consumption errors < 5% [30,31]. This shift improves comparability and promotes methodological rigor.
4. **Vehicle and Context Specific Operational Cycle Design**: Diesel and BEV cycle detection pipelines are increasingly tailored to the operating conditions and mechanical realities

of specific vehicle classes. Diesel studies differentiate loaders from haul trucks in both signal use and model architecture [3,17], while BEV research encodes domain-specific constraints such as gradient bins [32], shovel-resistance [30], or payload mass [31]. This marks a clear departure from generalized or automotive-style operational cycles.

4.2. Algorithm Families in Mining Cycle Detection

4.2.1. Mining-Specific Strengths and Weaknesses

Method families trade off data requirements, robustness, and complexity. Rule/threshold and change-point/segmentation approaches are simple and interpretable, often operating on a single hydraulic channel; they work well for coarse states (e.g., full/empty) but degrade on granular phases (loading vs. dumping), and are brittle to sensor noise or spikes [14,22,26]. On LHD hydraulics, segmentation exceeded 96% for full/empty driving yet only ≈67% for loading [12]. The HSMM family exploits sequence structure and durations and handles gaps via explicit "uncertain/offline" states, but can confuse spotting vs. dumping at low speeds [15]. Clustering (e.g., VBGMM) discovers regimes without labels but mixes operationally distinct transits unless features are carefully chosen [18]. With multi-sensor inputs (CAN speed/gear/RPM, IMU), trees/RF and CNN/LSTM deliver higher accuracy (typically 90–96%), with CNN/LSTM excelling on raw temporal signals and trees/RF competitive when strong features (e.g., spectra) are engineered. Common failure modes cluster at low-velocity states (idle/loading/unloading), where signatures overlap [3,11,17]. The studies are summarized in Table 4.

Table 4. Diesel methods: strengths, weaknesses, and best fit by signals/labels.

Family	Core Strengths	Core Weaknesses	Best Fit (Signals/Labels)	Typical Failure
Rule/Threshold	Simple; fast; transparent	Single-sensor brittleness	hydraulic pressure; clear gates; weak labels	Noise-triggered false cycles [14,22]
HMM/HSMM	Sequence + dwell modelling	Needs discretization; low-speed confusions	GPS/CAN; none/weak labels	Spotting vs. dump [15]
Change-point	Direct boundary finding	Smoothing-sensitive	HYD time-series; weak labels	Loading spikes [12,26]
Clustering	No labels required	Hard to name states	Vibration+speed; none labels	Mixed transit modes [18]
Trees/RF	Fast with features	Weaker on raw series	Multi-sensor CAN(+features); dense/weak	Feature drift/overfit [11]
CNN/LSTM	Learns temporal features	Needs labels; opaque	IMU/CAN raw; dense labels	Low-speed class overlap [3,11]
Hybrid	Rule guardrails + ML	Pipeline complexity	Multi-sensor + logic filters	Logic caps ML gains [4,10]

Let P_{dc} denote DC-bus power (positive under traction, negative under regenerative braking). During deceleration $P_{\mathrm{dc}} < 0$, while thermal/power and state-of-charge (SoC) limits clip $|P_{\mathrm{dc}}|$; steep grades skew energy-based features. Method families respond as follows: (i) HSMMs with signed emissions and minimum-dwell constraints reduce rapid label reversals at low speed; (ii) CNN/LSTM models learn patterns across $(P_{\mathrm{dc}}, \mathrm{SoC}, v)$ and capture energy-flow reversals without hand-crafted features; (iii) change-point detection on cumulative energy,

$$E(t) = \int_0^t P_{\mathrm{dc}}(\tau)\, d\tau,$$

is effective when the sign of energy flow persists; (iv) rules based only on $|P_{\mathrm{dc}}|$ degrade, but remain useful as guardrails (e.g., dump-approach gates).

Offline cycle-building often uses clustering, decision trees, Markov models, or hybrids on GPS/CAN data; evaluations emphasize representativeness rather than per-timestamp accuracy [31,32]. For buses, simple 1 Hz GPS segmentation is enough to separate coarse road-grade classes, whereas forklifts need higher-rate CAN and richer models to resolve fine-grained states. Markov designs can suffer from combinatorial growth, which asymmetric transition-probability matrices help mitigate.

4.2.2. Real-Time Readiness: Latency, Compute, and Robustness

Across model families, deployability depends on (i) latency (we distinguish sub-cycle <100 ms, intracycle 100–500 ms, intercycle 0.5–5 s), (ii) compute/energy (dominated by feature extraction and model size), and (iii) robustness to dropouts and noisy transitions. Simple rule pipelines and lightweight learners routinely meet *sub-cycle* budgets. To illustrate, in-house timing experiments on commodity CPUs (Intel i7-class, Python/ONNX runtime) show that a Random Forest operating on multivariate CAN windows executes in $\sim 3 \times 10^{-5}$ s per sample, while a compact 1-D CNN on the same inputs requires $\sim 4 \times 10^{-4}$ s, both with memory footprints in the single-digit MB range (these measurements are intended as order-of-magnitude guides; actual latencies depend on implementation, threading, and window length.). In practice, the bottleneck is often windowing or buffering rather than the classifier itself [11]. For high-rate channels (IMU/vibration, tens of Hz to kHz), the compute budget shifts toward feature extraction (e.g., FFTs) or compression on the edge device, while inference itself remains fast once features are materialized [18]. Models operating on coarse-rate streams or full segments (e.g., HSMMs over multi-second GPS) naturally align with *inter-cycle* updates [15]. Unsupervised clustering typically incurs a heavy offline training step but supports near-real-time assignment at run time. Robustness in practice relies on short-segment suppression, explicit *uncertain* states for gaps, and fallbacks to resilient ECU channels when fragile sensors degrade [14–16].

Several sources emphasize offline cycle construction rather than on-device inference; while they document training burdens (e.g., large transition matrices, kernel PCA), they provide limited direct latency or energy measurements for embedded deployment [31]. Where concrete timings are reported, they indicate the feasibility of sub-cycles for trees and small CNNs [11]; by contrast, sequence models and high-rate pipelines require careful windowing and, when necessary, compression or quantization to meet cycle targets.

Building on these run-time safeguards, the choice of edge platform sets the practical ceiling for latency, accuracy, and thermal headroom. In practice, options span always-on Neural Processing Units (NPUs) and high-performance Microcontroller Units (MCUs) (ultra-low power; robust to dust/vibration; limited model capacity) through edge Tensor Processing Unit (TPUs) and NPUs (balanced power/throughput; offloads feature work) to embedded Graphic Processing Units (GPUs) (max accuracy and flexibility; highest power and cooling needs). On MCU-class targets without Digital Signal Processor (DSP), windowing/Fast Fourier Transform can rival or exceed inference energy; on Edge-TPU/Jetson-class devices, inference dominates but requires stable power and heat rejection. For clarity, we use J/inference (pipeline-level) and kWh/shift (system-level); Table 5 maps representative model families to typical platforms and power envelopes.

Table 5. Quick reference: model family, representative edge platform(s), and typical power draw.

Model Family	Representative Edge Platform(s)	Typical Power
Rules/Thresholds	Always-on NPU (e.g., Syntiant-class [33]), High-perf MCU (e.g., STM32H7 [34])	<1 mW–0.7 W
HMM/HSMM	MCU (with DSP) or Edge NPU (for feature offload)	0.2–1 W
Trees/Random Forest	MCU (quantized features) [35] or Coral Edge TPU (post-feature)	0.5–3 W
1-D CNN (tiny)	Coral Edge TPU [36]/MCU+NPU combo	2–3 W
Bi-LSTM (tiny)	MCU (small hidden dims) or Edge NPU	0.5–2 W
CNN–LSTM (compact)	Coral Edge TPU/Jetson Orin Nano [37] (7/15 W modes)	3–8 W
Deep CNN/large hybrids	Jetson-class embedded GPU (Orin Nano/Xavier NX) [37,38]	7–15 W

Notes. Ranges are typical device-level draws under inference load and exclude peripherals (fans, radios, storage). Select platform power modes (e.g., 7 W vs. 15 W) to respect thermal limits; report both J/inference and kWh/shift.

4.2.3. Explainability Across Method Families

Method families in the corpus span from intrinsically transparent to largely opaque models. Rule-based pipelines can be directly audited by inspecting explicit thresholds and hysteresis (e.g., engine speed or brake pressure used to flag unloading) and checking them against known physics [14]. Probabilistic sequencers such as HMM/HSMM expose transition and emission matrices that can be reviewed for physical plausibility (e.g., prohibiting "reversing" observations during "loading"), enabling "what-if" interrogation of dwell and transition behavior [15]. Change-point detectors provide objective cost/penalty curves that justify detected boundaries and allow post-hoc scrutiny when segments are merged or split [16]. For tree-based ensembles, per-feature importance and path traces indicate which signals drive predictions [17]. By contrast, *deep sequence models* (CNN/LSTM) in the included mining studies were reported without saliency, occlusion, SHapley Additive exPlanations (SHAP) [39], or prototype analyses, leaving them effectively black-box in this literature [3,11,17,19].

Cross-domain BEV sources reinforce the same gradient of transparency. Markov-chain cycle builders publish transition probability matrices that visibly encode state-to-state logic, and manually constructed decision trees provide an explicit flow of splits (e.g., direction, gear, load mass) that practitioners can audit step by step [31]. Several BEV studies also add an energy-consistency check—comparing simulated battery SoC or power trajectories from the derived cycle to field traces—which functions as a domain-specific sanity check on mode assignments [31]. Taken together, most non-deep learning approaches arrive with built-in artifacts (rules, matrices, cost curves, feature importance), whereas deep nets will require adding post-hoc tools before they can be used in production.

Beyond noting that explainability differs across method families, it is useful to define a minimum set of artifacts that practitioners should routinely produce. These artifacts provide transparency for both engineers and operators by making classification logic, error modes, and trade-offs explicit. Table 6 summarizes, by model family, the key explainability outputs expected as part of a deployment "acceptance package".

Table 6. Recommended explainability artifacts by model family.

Model Family	Explainability Artifacts
Rule/Threshold	Explicit rule set; decision thresholds; coverage of fallback/uncertain states.
HMM/HSMM	Transition matrix; state-duration distributions; visualization of typical vs. anomalous paths.
Tree/Random Forest	Feature-importance ranking; SHAP summary plots; partial-dependence curves for top features.
Shallow neural nets (MLP, simple Recurrent Neural Networks)	Confusion and cost curves; sensitivity analysis to key input features.
CNN/LSTM hybrids	Saliency maps (Grad-CAM [40] or Integrated Gradients [41]) on input windows; per-channel attribution scores; confusion and transition matrices.
Unsupervised/Clustering	Cluster centroids/prototypes; distance distributions; visualization of borderline samples; mapping between clusters and cycle labels.

4.2.4. Task–Method Fit for Operational Objectives

Operational objectives emphasize different properties; matching families to task constraints avoids false comparisons across metrics/splits. The studies are summarized in Table 7.

Table 7. Task-oriented method fits across families.

Operational Task	Key Constraints	Well-Suited Families (Why)	Deployment Key Performance Indicators (KPIs)
Predictive maintenance	Rare-event precision; auditability; stable states	Hybrid rules + HSMM (dwell/transition logic, explainable); Trees/RF w/engineered features (feature importances)	Per-fault precision/recall; false alarms/100 h; time-to-detect
Energy optimization	Grade/regen awareness; segment energy	CNN/LSTM/BiLSTM over (P_{dc}, SoC, v); HSMM with signed emissions (captures regen dwell)	kWh/segment; grade-stratified Mean Absolute Error; latency
Shift KPIs/cycle counts	Low latency; robustness; few labels	Lightweight HMM/HSMM; rules with hysteresis; tiny CNNs	Cycle-count error; per-cycle F1; ms-latency

4.3. Research Gaps and Limitations

Several significant limitations remain unaddressed. First, there is a lack of publicly available data sets for mobile mining machinery (specifically, those collected for cycle detection), which is in sharp contrast to other industries; the automotive industry has numerous open datasets and standard benchmarks, such as the Fleet DNA for Duty Cycles dataset [42]. This shortage hampers reproducibility, prevents meaningful comparative analyses, and limits the engagement of the larger research community.

The bottleneck is not a lack of sensors—it is access and standardization. Mobile mining telemetry is fragmented across Original Equipment Manufacturer (OEM) portals and site systems, with unclear data ownership and proprietary signal dictionaries. Consequently, researchers rarely obtain shareable logs with ground truth. A practical remedy is a neutral, OEM agnostic corpus: raw, second-by-second time-series from LHDs, trucks, and jumbos (powertrain, hydraulics, brake/steer, bucket, payload, and SoC/voltage/temperature for BEVs), normalized to a common schema; per-timestamp operating mode labels; cycle boundary annotations; and per-cycle targets (e.g., tons moved, energy used/regenerated, charge/swap events). Release should include a cycle definition protocol per machine type (what events start/end a cycle, and modes make up a cycle) and a de-identification policy that obscures the manufacturer or site from which the data originated. The Global Mining Guidelines Group's open-data/interoperability work can guide the schema [43], but the dataset itself needs to be public.

Two automotive industry practices could translate cleanly to mining: (1) rich, open, labelled datasets (e.g., Fleet DNA for duty cycles; nuScenes, World-Harmonized Light-Duty Vehicles Test Procedure for transparent benchmarks [44,45]), and (2) public leaderboards with fixed splits and baseline implementations so methods are directly comparable. For underground mining, the benchmark should be time-aligned (event-anchored) rather than dependent on location, because the Global Navigation Satellite System is unavailable; camera/LiDAR streams can be helpful where feasible, but not required given dust, low light, and visual aliasing in underground environments. Primary tasks fit the operational cycle domain: mode labelling, cycle boundary detection, per-cycle energy/tonnage estimation, and BEV-specific metrics (e.g., regen ratio, SoC drop per cycle). To encourage generalization, include cross-site, multi-equipment splits and standardized metrics (e.g., F1-scores for modes/boundaries; Mean Absolute Percentage Error/Root Mean Squared Error for energy and SoC). Progress in mining cycle detection does not hinge on autonomous-vehicle style three-dimensional perception; it needs trustworthy, labelled time series, and a common scoreboard.

Inconsistencies in metrics and evaluation practices also present significant hurdles. **No standardized operational cycles exist for different classes of mining vehicles, and even among identical machine types, research diverges in defining precise operational states and transition points.** For example, when defining LHD operational cycles, ambiguity often arises around transitions between loading and hauling modes (i.e., does hauling begin when an LHD starts moving away from a muck pile?). To strengthen methodological rigor,

future benchmarks should also include a **standardized ablation protocol**: a preprocessing ablation suite (resampling rates, window sizes/strides, filter families and cutoffs, feature normalization) crossed with model baselines (rule-based, HMM/HSMM, tree-ensembles, CNN/LSTM). Researchers should report both per-segment and per-cycle metrics under identical splits (e.g., leave-asset and leave-mine) to enable fair attribution of gains.

From a modelling perspective, the potential of semi-supervised learning approaches remains significantly underutilized, despite their promise in addressing label scarcity and leveraging large volumes of unlabeled operational data. Moreover, the application of explainable machine learning techniques is notably absent from existing research; methods such as SHAP values for interpretability have not been explored within published cycle detection studies, representing a critical research gap that could otherwise improve practical adoption of current models.

Finally, substantial deployment challenges persist, marked by a notable gap between research demonstrations and production-ready, real-time cycle detection systems (only 3 out of the 20 cycle detection papers found through the search could be considered to be evaluated in real time). Addressing real-world constraints, such as computational resource limitations, sensor noise, and latency requirements, is essential to transition promising research results into robust and reliable tools that mining operators can practically adopt and integrate into existing fleet management and maintenance workflows.

Limitations of the review process: Beyond the field-level gaps identified above, the review process itself carries limitations. Only one bibliographic database (The Lens) was searched, which may have excluded relevant studies indexed elsewhere. The review was limited to English-language publications, and all screening and data extraction were carried out by a single reviewer, which introduces potential selection bias. To mitigate these risks, backward and forward citation chasing of sentinel papers and expert hand-searching of BEV exemplars were performed. However, it remains possible that some relevant studies were missed.

4.4. Future Research Opportunities

The primary barrier limiting progress in the detection of the mobile mining equipment operational cycle is the absence of open data sets, preventing meaningful comparison between studies. Developing and releasing a mining cycle identification benchmark data set (similar in style to the Turbofan Engine Degradation Simulation Data Set provided by NASA [46]), pairing annotated CAN and inertial streams with clear metrics, would enable consistent evaluations and facilitate direct comparison of methodologies.

Expanding research into unsupervised, semi-supervised, and self-supervised learning techniques (similar to the push seen in driving cycle literature for passenger vehicles) presents a promising opportunity. For instance, adopting techniques inspired by natural language processing (such as the Double Articulation Analyzer method [47]) could help better identify and characterize operational modes without extensive manual labelling.

Finally, there is considerable value in moving toward deployment-focused research, involving models tested in real-life operational settings. Future work could explore reinforcement learning-based post-processing techniques [48] to continuously refine the models initially trained offline, as well as the development and validation of real-time data pipelines. Incorporating lightweight, latency-aware model architectures, such as those demonstrated in Tiny-scale Machine Learning anomaly detection applications [1], could significantly enhance model suitability by allowing models to be deployed onto mobile mining equipment.

5. Conclusions

This review charted the evolution of operational cycle detection for mobile mining equipment, tracing the field's shift from single-sensor thresholds to multivariate, machine learning-based pipelines for diesel fleets, and examining the nascent literature on battery-electric mining vehicles. In doing so, the paper fulfills its stated objectives: (i) synthesizing and contrasting diesel and BEV methodologies, and (ii) exposing critical gaps—including the scarcity of open datasets, limited use of semi-/self-supervised learning, and the need for real-time, edge-ready solutions. Addressing these gaps will accelerate robust, easily deployed cycle detection systems for next-generation mining fleets.

Supplementary Materials: The following supporting information can be downloaded at https://www.mdpi.com/article/10.3390/eng6100279/s1: S1: Full Lens search string (executed 27 June 2025). Table S1: Title/abstract screening—exclusion reasons and counts. Table S2: Full-text assessment—exclusion reasons and counts. Table S3: Funding and conflict-of-interest statements (included + exemplar studies).

Author Contributions: Conceptualization, A.M.d.C.; Methodology, A.M.d.C.; software, A.M.d.C.; validation, A.M.d.C., M.T. and M.C.L.; formal analysis, A.M.d.C.; investigation, A.M.d.C.; resources, M.C.L.; data curation, A.M.d.C.; writing—original draft preparation, A.M.d.C.; writing—review and editing, A.M.d.C., M.T. and M.C.L.; visualization, A.M.d.C.; supervision, M.T. and M.C.L.; project administration, A.M.d.C. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding. As there was no external funding, there was no funder involvement in the design, data collection, analysis, decision to publish, or manuscript preparation. Funding and conflict-of-interest declarations for each included study are provided in Supplementary Materials (S3).

Data Availability Statement: No new data were created or analyzed in this study. Data sharing is not applicable to this article.

Acknowledgments: The authors gratefully acknowledge support from Natural Resources Canada—CanmetMINING and the Ontario Graduate Scholarship (OGS). During the preparation of this manuscript, the authors used ChatGPT (OpenAI, models: GPT-4o, GPT-4.5, and o3) to assist with study design and the initial drafting of text, and employed NotebookLM (Google, July 2025 release) to aid in the exploratory analysis and interpretation of literature-derived data. The authors have reviewed and edited all AI-generated output and take full responsibility for the content of this publication.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

BEV	Battery-Electric Vehicle
CAN	Controller Area Network (vehicle data bus)
CNN	Convolutional Neural Network
DBSCAN	Density-Based Spatial Clustering of Applications with Noise
ECU	Electronic Control Unit
GPS	Global Positioning System
KPI	Key Performance Indicator
VBGMM	Variational Bayesian Gaussian Mixture Model
HMM	Hidden Markov Model
HSMM	Hidden Semi-Markov Model
IMU	Inertial Measurement Unit
LHD	Load–Haul–Dump (vehicle)

EFS	Exhaustive Feature Selection
WMA	Weighted Moving Average
EMA	Exponential Moving Average
DEMA	Double-Exponential Moving Average
HMA	Hull Moving Average
ALMA	Arnaud Legoux Moving Average
GPU	Graphics Processing Unit
TPU	Tensor Processing Unit;
NPU	Neural Processing Unit
DSP	Digital Signal Processor
LOWESS	Locally Weighted Scatterplot Smoothing
LSTM	Long Short-Term Memory network
Bi-LSTM	Bidirectional Long Short-Term Memory network
SPEED	Vehicle Speed
GEAR/SELGEAR	Selected Gear
GA	Genetic Algorithm
OEM	Original Equipment Manufacturer
PCA	Principal Component Analysis
PRISMA-ScR	Preferred Reporting Items for Systematic Reviews and Meta-Analyses–Scoping Review extension
RPM	Revolutions Per Minute
SVM	Support Vector Machine
SoC	State of Charge (battery)
RF	Random Forest
TPM	Transition Probability Matrix

References

1. Fuente, R.d.l.; Radrigan, L.; Morales, A.S. Enhancing Predictive Maintenance in Mining Mobile Machinery through a TinyML-enabled Hierarchical Inference Network. *arXiv* **2024**, arXiv:2411.07168. [CrossRef]
2. Timusk, M.; Lipsett, M.G.; McBain, J.; Mechefske, C.K. Automated Operating Mode Classification for Online Monitoring Systems. *J. Vib. Acoust.* **2009**, *131*, 041003. [CrossRef]
3. Skoczylas, A.; Rot, A.; Stefaniak, P.; Śliwiński, P. Haulage Cycles Identification for Wheeled Transport in Underground Mine Using Neural Networks. *Sensors* **2023**, *23*, 1331. [CrossRef]
4. Karimi, G.; Masih-Tehrani, M.; Pourbafarani, Z. Development of Wheel Loader Duty Cycle Using Hybrid Markov Chain and Genetic Algorithm. *SAE Int. J. Commer. Veh.* **2021**, *15*, 51–64. [CrossRef]
5. Xiao, W.; Liu, M.; Chen, X. Research Status and Development Trend of Underground Intelligent Load-Haul-Dump Vehicle—A Comprehensive Review. *Appl. Sci.* **2022**, *12*, 9290. [CrossRef]
6. Long, M.; Schafrik, S.; Kolapo, P.; Agioutantis, Z.; Sottile, J. Equipment and Operations Automation in Mining: A Review. *Machines* **2024**, *12*, 713. [CrossRef]
7. Epiroc AB. Epiroc Wins Large Order for Battery-Electric Vehicles for Canadian Gold and Copper Mine. Press Release, 13 March 2025. Available online: https://www.epirocgroup.com/en/media/corporate-press-releases/2025/20250313-epiroc-wins-large-order-for-battery-electric-vehicles-for-canadian-gold-and-copper-mine (accessed on 15 October 2025).
8. Beltrami, D.; Iora, P.; Tribioli, L.; Uberti, S. Electrification of Compact Off-Highway Vehicles—Overview of the Current State of the Art and Trends. *Energies* **2021**, *14*, 5565. [CrossRef]
9. Ouzzani, M.; Hammady, H.; Fedorowicz, Z.; Elmagarmid, A. Rayyan—A web and mobile app for systematic reviews. *Syst. Rev.* **2016**, *5*, 210. [CrossRef]
10. Gawelski, D.; Jachnik, B.; Stefaniak, P.; Skoczylas, A. ICCCI (CCIS Volume)-Haul Truck Cycle Identification Using Support Vector Machine and DBSCAN Models. In *Communications in Computer and Information Science*; Springer: Cham, Switzerland, 2020; pp. 338–350. . [CrossRef]
11. Jakobsson, E.; Frisk, E.; Krysander, M.; Pettersson, R. Automated Usage Characterization of Mining Vehicles For Life Time Prediction. *IFAC-PapersOnLine* **2020**, *53*, 11950–11955. [CrossRef]
12. Koperska, W.; Skoczylas, A.; Stefaniak, P. ICCCI (CCIS Volume)—A Simple Method of the Haulage Cycles Detection for LHD Machine. In *Communications in Computer and Information Science*; Springer: Cham, Switzerland, 2020; pp. 326–337. [CrossRef]

13. Kozłowski, T.; Michalak, A.; Zimroz, R.; Sliwinski, P.; Kaniewski, T. An Algorithm for Determining Unloading Cycles in Dump Truck. *IOP Conf. Ser. Earth Environ. Sci.* **2019**, *362*, 012111. [CrossRef]
14. Krot, P.; Sliwinski, P.; Zimroz, R.; Gomolla, N. The identification of operational cycles in the monitoring systems of underground vehicles. *Measurement* **2020**, *151*, 107111. [CrossRef]
15. Markham, G.; Seiler, K.M.; Balamurali, M.; Hill, A.J. Load-Haul Cycle Segmentation with Hidden Semi-Markov Models-2022 IEEE 18th International Conference on Automation Science and Engineering (CASE). In Proceedings of the 2022 IEEE 18th International Conference on Automation Science and Engineering (CASE), Mexico City, Mexico, 20–24 August 2022; IEEE: Piscataway, NJ, USA, 2022; pp. 447–454. [CrossRef]
16. Polak, M. Identification of Loading Process Based on Hydraulic Pressure Signal-Sgem International Multidisciplinary Scientific GeoConference EXPO Proceedings. In Proceedings of the SGEM International Multidisciplinary Scientific GeoConference EXPO Proceedings, Albena, Bulgaria, 30 June–6 July 2016. [CrossRef]
17. Qi, Z.; Gu, Q.; Meng, Y.; Bai, G.; Ding, D. The load-haul-dump operation cycle recognition based on multi-sensor feature selection and bidirectional long short-term memory network. *Meas. Control* **2023**, *56*, 1523–1533. [CrossRef]
18. Saari, J.; Odelius, J. Detecting operation regimes using unsupervised clustering with infected group labelling to improve machine diagnostics and prognostics. *Oper. Res. Perspect.* **2018**, *5*, 232–244. [CrossRef]
19. Skoczylas, A.; Stefaniak, P.; Anufriiev, S.; Koperska, W. Deep Learning for Ore Haulage Monitoring: Vibrational Analysis Using a VGG16 Network. *IEEE Access* **2025**, *13*, 33138–33147. [CrossRef]
20. Śliwiński, P.; Andrzejewski, M.; Kaniewski, T.; Hebda-Sobkowicz, J.; Zimroz, R. *Mining Goes Digital-Selection of Variables Acquired by the On-Board Monitoring System to Determine Operational Cycles for Haul Truck Vehicle*; CRC Press: Boca Raton, FL, USA, 2019; pp. 525–533. [CrossRef]
21. Stefaniak, P.; Zimroz, R.; Obuchowski, J.; Sliwinski, P.; Andrzejewski, M. An Effectiveness Indicator for a Mining Loader Based on the Pressure Signal Measured at a Bucket's Hydraulic Cylinder. *Procedia Earth Planet. Sci.* **2015**, *15*, 797–805. [CrossRef]
22. Wodecki, J.; Michalak, A.; Stefaniak, P. Review of smoothing methods for enhancement of noisy data from heavy-duty LHD mining machines. *E3S Web Conf.* **2018**, *29*, 00011. [CrossRef]
23. Wodecki, J.; Góralczyk, M.; Krot, P.; Ziętek, B.; Szrek, J.; Worsa-Kozak, M.; Zimroz, R.; Śliwiński, P.; Czajkowski, A. Process Monitoring in Heavy Duty Drilling Rigs—Data Acquisition System and Cycle Identification Algorithms. *Energies* **2020**, *13*, 6748. [CrossRef]
24. Wyłomańska, A.; Zimroz, R. Signal segmentation for operational regimes detection of heavy duty mining mobile machines—A statistical approach. *Diagnostyka* **2014**, *55*, 33–42.
25. Zhang, X.; Wang, S.; Zhang, W. Research of Driving Cycle Construction for Electric Drive Mining Truck Based on Travelling Analysis Method. In Proceedings of the 2018 Chinese Intelligent Systems Conference, Wenzhou, China, 13–14 October 2018; Springer: Singapore, 2018; pp. 117–126. [CrossRef]
26. Stachowiak, M.; Koperska, W.; Jachnik, B.; Stefaniak, P. Tracking Operation Regimes of the LHD Machines Based on Hydraulic Signal Signal Smoothing Issues-Sgem International Multidisciplinary Scientific Geoconference Expo Proceedings. In Proceedings of the SGEM International Multidisciplinary Scientific GeoConference EXPO Proceedings, Albena, Bulgaria, 18–24 August 2020; Curran Associates, Inc.: New York, NY, USA, 2020; Volume 20, pp. 383–394. [CrossRef]
27. Lewis, D. A Machine Learning Approach to Underground Load Haul Dump Cycle Segmentation. Master's Thesis, School of Mechanical and Mining Engineering, The University of Queensland, Brisbane, Australia, 2018. [CrossRef]
28. Simonyan, K.; Zisserman, A. Very Deep Convolutional Networks for Large-Scale Image Recognition. *arXiv* **2015**, arXiv:1409.1556. [CrossRef]
29. Ye, Y.; Zhao, X.; Zhang, J. Driving cycle electrification and comparison. *Transp. Res. Part D Transp. Environ.* **2023**, *123*, 103900. [CrossRef]
30. Ren, H.; Xu, M.; Lin, T.; Chen, Q.; Cai, S.; Guo, T. Working cycle condition construction for electric wheel loader based on principal component analysis and cluster analysis. *Proc. Inst. Mech. Eng. Part D J. Automob. Eng.* **2024**, *238*, 2905–2919. [CrossRef]
31. Tong, Z.; Guan, S. Developing high-precision battery electric forklift driving cycle with variable cargo weight. *Transp. Res. Part D Transp. Environ.* **2024**, *136*, 104443. [CrossRef]
32. Tong, H.Y.; Ng, K.W. Developing electric bus driving cycles with significant road gradient changes: A case study in Hong Kong. *Sustain. Cities Soc.* **2023**, *98*, 104819. [CrossRef]
33. Syntiant Corporation. NDP120 Neural Decision Processor—Product Brief, 2023. Always-on Audio/Sensor NPU for Ultra-Low-Power Applications. Available online: https://static1.squarespace.com/static/6488b0b8150a045d2d112999/t/650b29219d8c3f0e27f4651e/1695230242406/Syntiant-Product_Brief_NDP120.pdf (accessed on 24 September 2025).
34. Emcraft Systems. Theoretically Calculated Max Power Consumption of the STM32H7 System-On-Module, 2024. Derives 220 mA @ 3.3 V (0.73 W) Upper-Bound Run Current from ST Datasheets. Available online: https://emcraft.com/som/stm32h7/stm32h7-power-consumption (accessed on 24 September 2025).

35. STMicroelectronics. STM32H742xI/G, STM32H743xI/G—32-bit Arm® Cortex®-M7 MCUs (Datasheet DS12110), 2023. Series Datasheet with Electrical Characteristics and Current Consumption Tables. Available online: https://www.syntiant.com/ndp120 (accessed on 24 September 2025).
36. Google LLC. Edge TPU Performance Benchmarks, 2025. States 4 TOPS at 0.5 W/TOPS (2 TOPS/W). Available online: https://gweb-coral-full.uc.r.appspot.com/docs/edgetpu/benchmarks/ (accessed on 24 September 2025).
37. NVIDIA Corporation. Jetson Orin Nano Series Data Sheet (DS-11105-001 v1.1), 2023. Module Datasheet; Orin Nano Series with 7-W and 15-W Reference Modes. Available online: https://www.mouser.com/pdfDocs/Jetson_Orin_Nano_Series_DS-11105-001_v11.pdf (accessed on 24 September 2025).
38. NVIDIA Corporation. Jetson Linux Developer Guide: Power Management for Jetson Xavier NX Series and Jetson AGX Xavier Series, 2024. Lists Supported 10-W, 15-W, and 30-W power budgets for Xavier NX. Available online: https://docs.nvidia.com/jetson/archives/r34.1/DeveloperGuide/index.html#page/Tegra%20Linux%20Driver%20Package%20Development%20Guide/power_management_jetson_xavier.html (accessed on 24 September 2025).
39. Lundberg, S.; Lee, S.I. A Unified Approach to Interpreting Model Predictions. *arXiv* **2017**, arXiv:1705.07874. .. [CrossRef]
40. Selvaraju, R.R.; Cogswell, M.; Das, A.; Vedantam, R.; Parikh, D.; Batra, D. Grad-CAM: Visual Explanations from Deep Networks via Gradient-based Localization. In Proceedings of the IEEE International Conference on Computer Vision (ICCV), Venice, Italy, 22–29 October 2017; pp. 618–626.
41. Sundararajan, M.; Taly, A.; Yan, Q. Axiomatic Attribution for Deep Networks. In Proceedings of the 34th International Conference on Machine Learning (ICML), Sydney, Australia, 6–11 August 2017; PMLR: Birmingham, UK, 2017; pp. 3319–3328.
42. Walkowicz, K.; Kelly, K.; Duran, A. *Fleet DNA*; National Renewable Energy Laboratory: Golden, CO, USA, 2014.
43. Global Mining Guidelines Group. Mobile Equipment Open Data: Report on Industry Challenges and Next Steps. Technical Report, Global Mining Guidelines Group. 2021. Available online: https://gmggroup.org/wp-content/uploads/2024/01/2021-03-17_Mobile-Equipment-Open-Data-Report-on-Industry-Challenges-and-Next-Steps.pdf (accessed on 20 September 2025).
44. Caesar, H.; Bankiti, V.; Lang, A.H.; Vora, S.; Liong, V.E.; Xu, Q.; Krishnan, A.; Pan, Y.; Baldan, G.; Beijbom, O. nuScenes: A multimodal dataset for autonomous driving. *arXiv* **2019**, arXiv:1903.11027. [CrossRef]
45. United Nations Economic Commission for Europe. UN Regulation No. 154—Worldwide harmonized Light Vehicles Test Procedure (WLTP). 2021. Available online: https://unece.org/transport/documents/2021/02/standards/un-regulation-no-154-worldwide-harmonized-light-vehicles-test (accessed on 27 September 2025).
46. Saxena, A.; Goebel, K. Turbofan Engine Degradation Simulation Data Set. NASA Prognostics Data Repository, NASA Ames Research Center, Moffett Field, CA, USA. 2008. Available online: https://www.nasa.gov/intelligent-systems-division/discovery-and-systems-health/pcoe/pcoe-data-set-repository/ (accessed on 5 August 2025).
47. Takenaka, K.; Bando, T.; Nagasaka, S.; Taniguchi, T.; Hitomi, K. Contextual scene segmentation of driving behavior based on double articulation analyzer. In Proceedings of the 2012 IEEE/RSJ International Conference on Intelligent Robots and Systems, Vilamoura-Algarve, Portugal, 7–12 October 2012; pp. 4847–4852. [CrossRef]
48. Radac, M.B.; Chirla, D.P. Near real-time online reinforcement learning with synchronous or asynchronous updates. *Sci. Rep.* **2025**, *15*, 17158. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.